What Computers Mean for Man and Society

Herbert A. Simon

Energy and information are two basic currencies of organic and social systems. A new technology that alters the terms on which one or the other of these is available to a system can work on it the most profound changes. At the core of the Industrial Revolution, which began nearly three centuries ago, lay the substitution of mechanical energy for the energy of man and animal. It was this revolution that changed a rural subsistence society into an urban affluent one and touched off a chain of technological innovations that transformed not only production but also transportation, communication, warfare, the size of human populations, and the natural environment.

It is easy, by hindsight, to see how inexorably these changes followed one another, how "natural" a consequence, for example, suburbia was of cheap, privately owned transportation. It is a different question whether foresight could have predicted these chains of events or have aided in averting some of their more undesirable outcomes. The problem is not that prophets were lacking—they have been in good supply at almost all times and places. Quite the contrary, almost everything that has happened, and its opposite, has been prophesied. The problem has always been to pick and choose among the embarrassing riches of alternative projected futures; and in this, human societies have not demonstrated any large foresight. Most often we have been constrained to anticipate events just a few years before their occurrence, or even while they are happening, and to try to deal with them, as best we can, as they are engulfing us.

We are now in the early stages of a revolution in processing information that shows every sign of being as fundamental as the earlier energy revolution. Perhaps we should call it the Third Information Revolution. (The first produced written language, and the second, the printed book.) This third revolution, which began more than a century ago, includes the computer but many other things as well. The technology of information comprises a vast range of processes for storing information, for copying it, for transmitting it from one place to another, for displaying it, and for transforming it.

Photography, the moving picture, and television gave us, in the course of a century, a whole new technology for storing and displaying pictorial information. Telegraphy, the telephone, the phonograph, and radio did the same for storing and transmitting auditory information. Among all of these techniques, however, the computer is unique in its capacity for manipulating and transforming information and hence in carrying out, automatically and without human intervention, functions that had previously been performable only by the human brain.

As with the energy revolution, the consequences of the information revolution spread out in many directions. First, there are the economic consequences that follow on any innovation that increases human productivity. As we shall see, these are perhaps the easiest effects of technological change to predict. Second, there are consequences for the nature of work and of leisure—for the quality of life. Third, the computer may have special consequences for privacy and individual liberty. Fourth, there are consequences for man's view of himself, for his picture of the universe and of his place and goals in it. In each of these directions, the immediate consequences are, of course, the most readily perceived. (It was not hard to foresee that Newcomen's and Watt's engines would change the economics of mining in deep pits.) It is far more difficult to predict what indirect chains of effects these initial impacts will set off, for example, the chain that reaches from the steam engine through the internal-combustion engine to the automobile and the suburb.

Prediction is easier if we do not try to forecast in detail the time path of events and the exact dates on which particular developments are going to occur, but to focus, instead, upon the steady state toward which the system is tending (*1*). Of course, we are not so much interested in what is going to happen in some vague and indefinite future as we are in what the next generation or two holds for us. Hence, a generation is the time span with which I shall be concerned.

My discussion will be divided into five parts, the last four corresponding to domains of prediction: economics, the nature of work and leisure, social consequences, and how men and women view themselves. These essays in prediction need to be preceded, however, by some analysis of the computer itself, and particularly its capabilities and potential in the area that is usually called artificial intelligence. This subject is taken up in the next section (*2*).

Computer Capabilities

The computer is a device endowed with powers of utmost generality for processing symbols. It is remarkable not only for its capabilities but also for the simplicity of its underlying processes and organization. Of course, from a hardware standpoint it is not simple at all but is a highly sophisticated electronic machine. The simplicity appears at the level of the elementary information processes that the hardware enables it to perform, the organization for execution and control of those processes, and the programming languages in terms of which the control of its behavior is expressed. A computer can read symbols from an external source, output symbols to an external destination, store symbols in one or more memories, copy symbols, rearrange symbols and structures of symbols, and react to symbols conditionally—that is, follow one course of action or another, depending on what symbols it finds in memory or in its input devices. The most general symbol-manipulating system that has been defined, the so-called Turing machine, requires no broader capabilities than these. The important limits on the powers of a computer are limits on the sizes of its memories and the speed of its elementary processes, and not on the generality of those processes.

There is great dispute among experts as to what the generality of the computer implies for its ability to behave intelligently. There is also dispute as to whether the computer, when it is behav-

The author is professor of computer science and psychology at the Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

ing more or less intelligently, is using processes similar to those employed by an intelligent human being, or quite different processes. The views expressed here will reflect my own experience in research with computers and my interpretation of the scientific literature. First, no limits have been discovered to the potential scope of computer intelligence that are not also limits on human intelligence. Second, the elementary processes underlying human thinking are essentially the same as the computer's elementary information processes, although modern fast computers can execute these processes more rapidly than can the human brain (*3*). In the past, computer memories, even in large computers, have probably not been as capacious as human memory, but the scale of available computer memories is increasing rapidly, to the point where memory size may not be much longer an effective limit on the capacity of computers to match human performance. Any estimate of the potential of the computer in the near or distant future depends on one's agreement or disagreement with these assumptions.

One common objection to the beliefs just expressed is that "computers can only do what you program them to do." That is correct. The behavior of a computer at any specific moment is completely determined by the contents of its memory and the symbols that are input to it at that moment. This does not mean that the programmer must anticipate and prescribe in the program the precise course of its behavior. A program is not a scenario; it is a strategy of action, and what actions actually transpire depends on the successive states of the machine and its inputs at each stage of the process—neither of which need be envisioned in advance either by the programmer or by the machine. A problem-solving program applied to a particular puzzle situation does not prescribe all the steps to solve that puzzle; it prescribes a selective search strategy that, when followed, may lead the computer to discover a path to a solution. But selective search, under the guidance of strategies, is also the process that people use to solve puzzles.

Of course humans, through the processes called learning, can improve their strategies by experience and instruction. By the same token, computers can be, and to some extent have been, provided with programs (strategies) for improving their own strategies. Since a computer's programs are stored in the same memories as data, it is entirely possible for programs to modify themselves—that is, to learn.

Probably the most fundamental differences between today's computers and the human information-processing system have to do with the input organs that provide the interface between the system and its environment. Simulating the capabilities of human eyes and ears has proved to be a much more difficult task than simulating the thinking processes that go on in the central nervous system. Computer capabilities in both visual and auditory domains, and particularly the former, fall far short of human capabilities.

Over the past two decades a moderate amount of work has been carried on in the field usually called artificial intelligence to explore the potentialities of the computer that have been outlined above. Some of this research has been aimed at programming computers to do things which, if done by a person, would be regarded as intelligent. Another part of the research has been directed at simulating not only the human capabilities but also the processes that human beings use in exercising these capabilities. The considerable progress that has been made in understanding the nature both of artificial and of human intelligence has hardly begun to translate itself into applications, and has been reflected to only a small degree in the actual practical uses of computers. Artificial intelligence research has had an impact upon the search algorithms that are used to solve large combinatorial problems, it is on the verge of practical application in the realm of medical diagnosis, and it has had an important influence upon certain computer programming techniques (for example, list processing). But its main significance for practical affairs lies in the future.

How, then, have computers actually been used to date? At present, computers typically spend most of their time in two main kinds of tasks: carrying out large-scale engineering and scientific calculations and keeping the financial, production, and sales records of business firms and other organizations. Although precise statistics are not available, it would be safe to estimate that 95 percent of all computing power is allocated to such jobs. Now these tasks belong to the horseless-carriage phase of computer development. That is to say, they consist in doing things rapidly and automatically that were being done slowly and by hand (or by desk calculator) in the precomputer era.

Such uses of computers do not represent new functions but only new ways of performing old functions. Of course, by greatly lowering the cost of performing them, they encourage us to undertake them on a larger scale than before. The increased analytic power provided by computers has probably encouraged engineers to design more complex structures (for example, some of the very tall new office buildings that have gone up in New York and Chicago) than they would have attempted if their analytic aids were less powerful. Moreover, by permitting more sophisticated analyses to be carried out in the design process, they have also brought about significant cost reductions in the designs themselves. In the same way, the mechanization of business record-keeping processes has facilitated the introduction of improved controls over inventories and cash flows, with resulting savings in costs. Thus, the computer not only reduces the costs of the information-processing operations that it automates but also contributes to the productivity of the activities themselves.

The remaining 5 percent of computer uses are more sophisticated. Let us consider two different ways in which a computer can assist an engineer in designing electric motors. On the one hand, the engineer can design the motor using conventional procedures, then employ the computer to analyze the prospective operation of the design—the operating temperature, efficiency, and so on. On the other hand, the engineer can provide the computer with the specifications for the motor, leaving to the computer the task of synthesizing a suitable design. In the second, but not the first, case the computer, using various heuristic search procedures, actually discovers, decides upon, and evaluates a suitable design. In the same way, the role of the computer in managing inventories need not be limited to record-keeping. The computer program may itself determine (on the basis of usage) when items should be reordered and how large the orders should be. In these and many other situations, computers can provide not only the information on which decisions are made but can themselves make the decisions. Process-control computers, in automated or semiautomated manufacturing operations, play a similar role in decision-making. Their programs are decision strategies which, as the system's variables change from moment to moment, retain control over the ongoing process.

It is the capability of the computer for solving problems and making decisions that represents its real novelty and that poses the greatest difficulties in predicting its impact upon society. An enormous amount of research and developmental activity will have to be carried out before the full practical implications of this capability will be understood and available for use. In the single generation that modern computers have been in existence, enough basic research has been done to reveal some of the fundamental mechanisms. Although one can point to a number of applications of the computer as decision-maker that are already 20 or 25 years old, development and application on a substantial scale have barely begun.

Economic Effects of Computers

The direct economic effects of introducing computers as numerical calculators and decision-makers are like those of introducing any new form of capital that raises productivity and also improves the quality of the product. The computer (its hardware together with the associated system-programming costs) represents an investment in a capital-intensive, laborsaving device that has to justify itself, in competition with other possible forms of investment, through savings in clerical and other personnel costs together with the improvements it brings about in organizational decisions.

When the main motive of introducing the computer is to mechanize existing clerical operations—in the actuarial department of an insurance firm, say, or the accounting department of a manufacturing concern—then its main economic advantage stems from the reduction in clerical costs. When it is introduced to mechanize decision processes—engineering design, for example, or control of stock or cash inventories—then its direct effect shows up as some form of productivity increase in the organization's operations. In either case, there is nothing special about the computer that distinguishes it, in its economic effects, from any other capital investment. Any such investment can be expected to have a direct effect upon employment in the organizational components where it is introduced. When the accounting system is mechanized, fewer clerks and bookkeepers are needed, else there would be no economic motivation for mechanizing. Of course, if part of the motivation for the change is to improve the quality of the system's output, the operation may be expanded, and the net reduction in personnel may be smaller than would be estimated solely from the increase in efficiency. If there is sufficient elasticity of demand for the activity, personnel may actually increase.

The most important question, however, is what the reduction in personnel at the point of impact means for the total level of employment in the economy. Again, this is a general economic issue—of technological unemployment—that does not depend on any special properties of computers. They are simply one among the many laborsaving devices that have been appearing since the beginning of the Industrial Revolution (and before).

Both standard economic analysis and a large body of empirical evidence demonstrate that there is no relation, positive or negative, between the technological sophistication of an economy and the level of employment it maintains. From a systems standpoint, a cost reduction in any part of the system releases resources that can be employed to increase the output of goods and services elsewhere in the system. At any level of employment, from 0 to 100 percent, the total revenue received by wage earners and owners of capital and land as wages, interest, and rent is just sufficient to purchase the total bundle of goods and services that is produced. Economists sometimes disagree as to why economies do not always operate at or near full employment, but they are unanimous in agreeing that the reason is not that they produce more than they can consume. (Even Marxists agree with this proposition, although they argue that full employment cannot be maintained within the institutions of capitalism.)

An even stronger statement can be made about the systems effects of cost-saving technological innovations. We usually describe devices like computers (and most other machinery) as laborsaving because they require a lower ratio of labor to capital than the methods they displace. But if we measure savings relative to output, they are usually both laborsaving and capital-saving. That is to say, a smaller capital investment per passenger mile is required to transport people by jet plane than to transport them by ox cart. Similarly, a smaller capital investment per multiplication is required if a large modern computer is used to do the arithmetic than if it is done on a desk calculator or with pencil and paper. (Do not forget to include the capital cost of the desk at which the clerk sits and the heated or air-conditioned room in which he or she works.) Now it is easy to show, for economic equilibrium and under reasonable assumptions about the supply of capital, that the introduction of capital-intensive, cost-saving innovations will raise the level of real wages and increase the fraction of the total revenue that goes to wages. This prediction from economic theory is amply supported by the histories of the industrialized economies over the present century. As productivity has increased (mainly as a consequence of technological innovation), real wages have steadily risen, as has labor's share in the total national product (*4*).

Now the rate of technological change depends both upon the rate of discovery of new innovations and upon the availability of capital to turn them into bricks and steel (or wire and glass). In this process, computers compete with other forms of technology for the available capital. Hence, the process of computerization is simply a part, currently an important part, of the general process of technological change. It might be described, paraphrasing Clausewitz, as "a continuation of the Industrial Revolution by other means" (*5*).

In taking this very global and bird's-eye view of the economics of mechanization, we should not ignore the plight of the worker who is displaced by the computer. His plight is often genuine and serious, particularly if the economy as a whole is not operating near full employment, but even if it is. Society as a whole benefits from increased productivity, but often at the expense of imposing transient costs on a few people. But the sensible response to this problem is not to eschew the benefits of change; it is rather to take institutional steps to shift the burdens of the transition from the individual to society. Fortunately, our attitudes on these questions appear to be maturing somewhat, and our institutional practices improving, so that the widespread introduction of the computer into clerical operations over the past generation has not called forth any large-scale Ludditism. In fact, during the depression that we are currently experiencing, in contrast to some earlier ones, technology has not been accused as the villain.

Effects on the Nature of Work

We see that, so far as economic effects are concerned, the computer simply provides a particular path toward higher productivity through industrialization. Whatever benefits it produces, it produces in this way; whatever problems it creates, it creates as other capital-intensive innovations do. We must be careful, however, not to evaluate social change solely in terms of its impact on wages and employment. Of equal importance are the effects it may have on the workplace, and even on leisure. Today we frequently hear the claim that computers and automation dehumanize work and that dehumanization, in turn, causes alienation from work and society. These charges have been laid not only against contemporary developments in automation but against the whole process of industrialization. They were stated eloquently in the *Communist Manifesto* more than a century ago and by numerous social critics before and since. There has been a new surge of concern with the alienation issue in the past 10 years.

Three questions need to be asked about alienation. First, how much alienation is there—is there evidence that alienation has been increased by computers and automation or, for that matter, by other forms of industrialization and mechanization? Second, in what ways is the nature of work, and the satisfactions derivable from it, changed by automation of the workplace? Third, as automation eliminates certain kinds of jobs in favor of others, what are the net effects upon the profile of jobs in the economy—are the jobs that are eliminated, on balance, more or less satisfying than the new ones that are created to replace them?

Objective data on national trends in job satisfaction are available only for about the last 20 years. About 15 national surveys have been conducted since 1958 by professional polling organizations that included questions on job satisfaction. Although 20 years is a short time, it does cover almost the whole period of the introduction of computers; hence these data should help answer the question before us. The polls provide absolutely no evidence for a decrease in job satisfactions over this period. If alienation has been increased by automation, the increase somehow does not show up in answers by workers to questions about their attitudes toward their jobs (*6*). Notice that the polls do not show that workers are enthusiastic about their jobs, only that they do not seem to like them less today than they did in 1958.

Unfortunately, comparable data are not available to measure the longer trends in job satisfaction over the whole past two centuries or so of industrialization. Perhaps even if computers and automation do not intensify alienation, they confirm and complete a loss of satisfactions that was produced by the rise of the factory system. The answer to that question must be mainly speculative. Clayre, however, recently threw some interesting light on it by examining the attitudes toward work expressed in preindustrial folk literature and popular ballads (*7*). He finds few indications of a Golden Age in which work was generally regarded as pleasurable and satisfying. He concludes that, in general, daily work was the same burdensome necessity for peasants and craftsmen as it is for factory workers and clerks; and that life's satisfactions and pleasure were mainly sought, then as now, in leisure, not work.

Perhaps, however, we should not try to detect alienation in this indirect way but should look directly at the workplaces where computers have been introduced, in order to see how they have changed the nature of work and its environment. Sizable differences have been found in worker satisfaction among blue-collar workers in different kinds of factories, some of the important variables being job variety and worker control over the timing of work. It is not the case, however, that the most advanced forms of industrialization and automation produce the most tedious and restrictive jobs. On the contrary, those forms of work organization that appear to have been most alienating—typified by the auto assembly line or large-scale hand assembly operations—are declining in importance relative to other forms of mechanization. Blauner, for example, studied four industries in depth: printing, a traditional craft industry; textiles, a machine-tending industry; automobile assembly, highly mechanized with highly specialized jobs; and a highly automated continuous-process chemical manufacturing industry (*8*). He found few indications of alienation in printing and chemicals, considerably more in textiles, and most of all in automobile assembly. The industry that best typifies modern automation—chemicals—was substantially less alienating than the two that typify older kinds of mechanization.

If we look at office automation, we see that, here too, the kinds of jobs that are displaced tend to be those that are most repetitive and restricting. Whisler, who studied about 20 companies in the insurance industry, found that computerization had produced only small and conflicting changes in the nature of clerical and supervisory jobs (*9*). The new jobs placed greater demands on the employees for accuracy and reliability in performance, but they were not generally perceived as being significantly more or less pleasant or more or less boring than before. And, perhaps most important of all, whatever effects were produced were small effects. Automation and computerization do not appear to change the nature of work in a fundamental way.

Again we must look not just at immediate impact but at system effects. Factory and office automation are laborsaving technologies. The jobs they eliminate are mostly jobs that were already relatively routine. Therefore, when we look at the impact on the labor force as a whole, we expect to see automation bringing about an overall decrease in the percentage of persons engaged in routine work of these kinds. Correspondingly, there will be a larger percentage of employees than before in service occupations and probably also in technical occupations. The work-satisfaction studies discussed earlier show differences among occupational groups of precisely this kind. From these data it appears that if factory operatives and clerical workers decline as a fraction of the labor force, while service workers, sales personnel, and professional and technical workers increase, there will be a net increase in reported job satisfaction—unless, of course, a compensating shift takes place in aspirations, a possibility we must not dismiss.

On all counts, then, we must acquit the computer technology of the charges that it has been and will be a cause of widespread alienation from work. Empirically, we find no signs of a downward trend in work satisfactions, and when we look at the actual impact of automation upon the workplace and the work force, we find no reason why such a trend should be expected. On the contrary, the newer technologies may even have a modest humanizing effect on the nature of work. The notion of a Golden Age of work prior to the Industrial Revolution must also be dismissed as romanticism, unsupported by such evidence as has been examined.

Control and Privacy

The potential of computers for increasing the control of organizations or society over their members and for invading the privacy of those members has caused considerable concern. The issues are important but are too complex to be discussed in detail here. I shall therefore restrict myself to a few comments which will serve rather to illustrate this complexity than to provide definitive answers.

A first observation is that our concern here is for competitive aspects of society, the power of one individual or group relative to others. Technologies tend to be double-edged in competitive situations, particularly when they are available to both competitors. For example, the computerization of credit information about individuals facilitates the assembly of such information from many sources, and its indefinite retention and accessibility. On the other hand, it also facilitates auditing such information to determine its sources and reliability. With appropriate legal rules of the game, an automated system can provide more reliable information than a more primitive one and can be surrounded by more effective safeguards against abuse. Some of us might prefer, for good reasons or bad, not to have our credit checked at all. But if credit checking is a function that must be performed, a strong case can be made for making it more responsible by automating it, with appropriate provision for auditing its operation.

Similarly, much has been said of the potential for embezzlement in computerized accounting systems, and cases have occurred. Embezzlement, however, was known before computers, and the computer gives auditors as well as embezzlers powerful new weapons. It is not at all clear which way the balance has been tilted.

The privacy issue has been raised most insistently with respect to the creation and maintenance of longitudinal data files that assemble information about persons from a multitude of sources. Files of this kind would be highly valuable for many kinds of economic and social research, but they are bought at too high a price if they endanger human freedom or seriously enhance the opportunities of blackmailers. While such dangers should not be ignored, it should be noted that the lack of comprehensive data files has never been the limiting barrier to the suppression of human freedom. The Watergate criminals made extensive, if unskillful, use of electronics, but no computer played a role in their conspiracy. The Nazis operated with horrifying effectiveness and thoroughness without the benefits of any kind of mechanized data processing.

Making the computer the villain in the invasion of privacy or encroachment on civil liberties simply diverts attention from the real dangers. Computer data banks can and must be given the highest degree of protection from abuse. But we must be careful, also, that we do not employ such crude methods of protection as to deprive our society of important data it needs to understand its own social processes and to analyze its problems.

Man's View of Man

Perhaps the most important question of all about the computer is what it has done and will do to man's view of himself and his place in the universe. The most heated attacks on the computer are not focused on its possible economic effects, its presumed destruction of job satisfactions, or its threats to privacy and liberty, but upon the claim that it causes people to be viewed, and to view themselves, as "machines" (*10*).

To get at the real issues, we must first put aside one verbal confusion. All of us are familiar with a wide variety of machines, most of which predated the computer. Consequently, the word "machine" carries with it many connotations: of rigidity, of simplicity, of repetitive behavior, and so on. If we call anything a machine, we implicitly attribute these characteristics to it. Hence, if a computer is a machine, it must behave rigidly, simply, and repetitively. It follows that computers cannot be programmed to behave like human beings.

The fallacy in the argument, of course, lies in supposing that, because we have applied the term "machine" to computers, computers must behave like older forms of machines. But the central significance of the computer derives from the fact that it falsifies these earlier connotations. It can, in fact, be programmed to behave flexibly, in complex ways, and not repetitively at all. We must either get rid of the connotations of the term, or stop calling computers "machines."

There is a more fundamental question behind the verbal one. It is essentially the question that was raised by Darwinism, and by the Copernican revolution centuries earlier. The question is whether the dignity of man, his sense of worth and self-respect depends upon his being something special and unique in the universe. As I have said elsewhere (*2*, p. 27):

> The definition of man's uniqueness has always formed the kernel of his cosmological and ethical systems. With Copernicus and Galileo, he ceased to be the species located at the center of the universe, attended by sun and stars. With Darwin, he ceased to be the species created and specially endowed by God with soul and reason. With Freud, he ceased to be the species whose behavior was—potentially—governable by rational mind. As we begin to produce mechanisms that think and learn, he has ceased to be the species uniquely capable of complex, intelligent manipulation of his environment.

What the computer and the progress in artificial intelligence challenge is an ethic that rests on man's apartness from the rest of nature. An alternative ethic, of course, views man as a part of nature, governed by natural law, subject to the forces of gravity and the demands of his body. The debate about artificial intelligence and the simulation of man's thinking is, in considerable part, a confrontation of these two views of man's place in the universe. It is a new chapter in the vitalism-mechanism controversy.

Issues that are logically distinct sometimes become stuck together with the glue of emotion. Several such issues arise here:

To what extent can human behavior be simulated by computer?

In what areas of work and life should the computer be programmed to augment or replace human activities?

How far should we proceed to explore the human mind by psychological research that makes use of computer simulation?

The first of these three issues will only be settled, over the years, by the success or failure of research efforts in artificial intelligence and computer simulation. Whatever our beliefs about the ultimate limits of simulation, it is clear that the current state of the art has nowhere approached those limits.

The second question will be settled anew each year by a host of individual and public decisions based on the changing computer technology, the changing economics of computer applications, and our attention to the social consequences of those applications.

The answer to the third question depends upon our attitudes toward the myths of Pandora and Prometheus. One viewpoint is that knowledge can be dangerous—there are enough historical examples—and that the attempt to arrive at a full explanation of man's ability to think might be especially dangerous. A different point of view, closer to my own, is that knowledge is power to produce new outcomes, outcomes that were not previously attainable. To what extent these outcomes will be good or bad depends on the purposes they serve, and it is not easy, in advance, to predict the good and bad uses to which any particular technology will be put. Instead, we must look back over human history and try to assess whether, on balance, man's gradual emergence from a state of ignorance about the world and about himself has been something we should celebrate or regret. To believe that knowledge is to be preferred to ignorance is to believe that the human species is capable of progress and, on balance, has progressed over the centuries. Knowledge about the human mind can make an important contribution to that progress. It is a belief of this kind that persuades researchers in artificial intelligence that their endeavor is an important and exciting chapter in man's great intellectual adventure.

Summary

From an economic standpoint, the modern computer is simply the most recent of a long line of new technologies that increase productivity and cause a gradual shift from manufacturing to service employment. The empirical evidence provides no support for the claim sometimes made that the computer "mechanizes" and "dehumanizes" work. Perhaps the greatest significance of the computer lies in its impact on Man's view of himself. No longer accepting the geocentric view of the universe, he now begins to learn that mind, too, is a phenomenon of nature, explainable in terms of simple mechanisms. Thus the computer aids him to obey, for the first time, the ancient injunction, "Know thyself."

References and Notes

1. A few years ago, Newell and I erred in predicting that certain specific developments in artificial intelligence were going to take place "within 10 years." The fact that we were optimistic about the time scale has blinded a number of critics to the basic soundness of our characterization of the nature and directions of artificial intelligence. I shall try not to make the same mistake here of predicting that very specific things will occur at very specific times. See H. A. Simon and A. Newell, *Oper. Res.* **6**, 1 (1958).
2. For more detailed discussions of these topics, see H. A. Simon, *The New Science of Management Decision* (Prentice-Hall, Englewood Cliffs, N.J., ed. 3, 1977).
3. The position that computers can be programmed to simulate an indefinite range of human thinking processes is developed in detail, and a large body of supporting evidence is examined in A. Newell and H. A. Simon, *Human Problem Solving* (Prentice-Hall, Englewood Cliffs, N.J., 1972); and in J. R. Anderson and G. H. Bower, *Human Associative Memory* (Winston, Washington, D.C., 1973).
4. For a fuller discussion of this evidence, see (*2*, chap. 4).
5. K. von Clausewitz, *On War* (Modern Library, New York, 1943), p. 596.
6. The polls under discussion were conducted by the Survey Research Centers of the Universities of Michigan and California, the National Opinion Research Center, and the Gallup Poll. For a detailed analysis of these data, see R. P. Quinn and L. J. Shepard, *The 1972–73 Quality of Employment Survey* (Institute for Social Research, University of Michigan, Ann Arbor, 1974).
7. A. Clayre, *Work and Play* (Harper & Row, New York, 1974).
8. R. Blauner, *Alienation and Freedom: The Factory Worker and His Industry* (Univ. of Chicago Press, Chicago, 1964).
9. T. L. Whisler, *The Impact of Computers on Organizations* (Praeger, New York, 1970). I. R. Hoos [in *Automation in the Office* (Public Affairs Press, Washington, D.C., 1961)] reaches more pessimistic conclusions than Whisler about the impact of the computer, but she mainly observed transient effects at the time the new technology was being introduced, and not the longer-term effects after the changes had been digested.
10. Two books that attack artificial intelligence on this ground are H. L. Dreyfus, *What Computers Can't Do* (Harper & Row, New York, 1972) and J. Weizenbaum, *Computer Power and Human Reason* (Freeman, San Francisco, 1976). The two books have little in common except a shared antipathy against the "machine" view of human thinking, and an eloquent contempt for those who hold that view. Weizenbaum's book is the technically more competent of the two, with respect to its understanding of the current state of the computer programming art.

Trends in Computers and Computing: The Information Utility

Stuart E. Madnick

The complexity, interdependence, and rapidity of events in modern society have accelerated demands for more effective ways to store, process, and manage information. Advances in both computer hardware (electronics) and software (programming) have provided the technology that can make it possible to effectively address many of these demands. In this article I review the structure of classical computer-based information systems and then consider these advances and show how they fit into the evolution of the information utility.

Information Systems

The need for more effective information management for both decision-making and operational efficiency is being felt for many reasons, such as:

1) Scarce resources. As we have become more aware of the scarcity of natural resources and the disruptions and perturbations in their supply (such as cost increases, boycotts, and strikes), it has become necessary to be much more responsive in decision-making in both the public and the private sectors.

2) Service productivity. There has been a growing demand for human services in most of the industrialized world. More than 66 percent of the American work force is employed in providing services; this includes teachers, lawyers, accountants, bank tellers, mailmen, and office workers. As noted by Hollomon (*1*), "The U.S. has evolved into a service economy and productivity is hard to manage in service-oriented industry. . . . Information technology is the most important element in a service economy." Information systems can improve the management of complex human services and provide greater operational efficiency, for example, through electronic funds transfer, electronic mail, automated office equipment, and electronic filing.

3) Complexity of society. As society becomes more complex, it is increasingly difficult to effectively assess the subtle interdependence of many decisions. As a result, "optimal" decision to resolve one problem may, in fact, precipitate worse problems in other areas. The frequency and significance of these counter-intuitive reactions are increasing (*2*) and this has prompted efforts to develop effective decision-support systems (*3*).

The element common to all of these concerns is information—to make decisions that can minimize the negative impacts of limited resources, to effectively manage the distribution of human services, and to cope with the complexity of society. Recent advances in computer and communications electronics, system design, and management science provide an increased potential for developing highly effective and efficient information systems.

The author is an associate professor of management science at the Alfred P. Sloan School of Management, a cofounder of the Center for Information Systems Research, and an adjunct member of the Laboratory for Computer Science (formerly Project MAC), all of the Massachusetts Institute of Technology, Cambridge 02139. This article is based on an internal report prepared in cooperation with Dean Witter & Co., Inc.